

Fortran direct access IO with MPI? ...

[clune](#) 113 posts since

May 31, 2007 I am curious as to whether the following I/O strategy is generally portable and/or standard conforming. (Note - I am not proposing this as an ideal solution for I/O, but rather just want to know whether it should be expected to "work" on all systems.)

- 1) Open file with direct access
- 2) Different processes write to *distinct* records within the file

My memory is that such a strategy would not work on the Cray T3E unless the record lengths conformed with some filesystem (blocksize?) constraints. However, playing around with a test code seems to indicate that it will work with arbitrary record lengths on our local Linux cluster (GPFS) and my Mac laptop. My concern is that when multiple processes attempt to update a block on the disk there is a race condition that must be handled by the filesystem. Do all filesystems handle this race condition? (I.e. atomically update blocks) If not, which common filesystems do not? If yes, is there some place where this is documented?

My test code is below (runs on just 2 processors):

```
program main
implicit none

integer, parameter :: N = 62
integer, parameter :: M = 100
real(kind=8) :: x(N,M), Y(N,M)
integer :: i

include 'mpif.h'
integer :: ier, rank

open(10,file='foo.dat',status='unknown',form='unformatted',access='direct', recl=2*N)

call mpi_init(ier)
call mpi_comm_rank(MPI_COMM_WORLD, rank, ier)
call random_number(y)
x=y
call mpi_barrier(MPI_COMM_WORLD, ier)
do i = 1, M
write(10,rec= getRecordNumber(i,rank)) x(:,i)
end do

close(10)
call mpi_barrier(MPI_COMM_WORLD, ier)

open(10,file='foo.dat',status='old',form='unformatted',access='direct', recl=2*N)
```

Fortran direct access IO with MPI? ...

```
do i = 1, M
  read(10,rec=getRecordNumber(i,rank)) x(:,i)
  if (.not. all(x(:,i) == y(:,i))) print*, 'uh oh: ', i
end do

close(10)

call mpi_finalize(ier)

contains

integer function getRecordNumber(i, rank) result(recNum)
integer, intent(in) :: i
integer, intent(in) :: rank

recNum = 2*i - 1 + rank
end function getRecordNumber

end program main
```

Message was edited by: clune Tags: mpi, fortran, direct_access, io

[crpearso](#) 1 posts since

May 8, 2008 **1. Re: Fortran direct access IO with MPI?** Nov 14, 2008 3:00 PM

As a starting point, I've asked the systems folks for some reference docs specific to Discover.